

In the Claims

Claims remaining in the application are as follows:

1. (Currently Amended): A database system ~~adapted to execute for~~ executing a database application that transfers a logical object in multiple fragments, ~~the database system~~ comprising:

a main storage site;

a remote storage site ~~adapted to link that links to the main storage site and to receive and store~~ receives and stores mirror information stored in from the main storage site, the remote storage site including and comprises a storage and a cache sidefile divided into a plurality of array sidefile recordsets;

a main protocol executable on the main storage site ~~and adapted to transfer that~~ transfers the logical object in multiple fragments in combination with information indicative of logical object fragment commencement and completion in the multiple fragment database application transfer comprising: receiving an application request to write the logical object of a specified length to a specified virtualized storage address; converting the virtualized write address and resolving the transfer length to designate at least one physical address in at least one physical storage device for transferring the logical object in fragments; sending a first control message to the at least one physical storage device identifying the start of a logical object to be held in a remote mirror cache for destaging; and sending a second control message identifying the end of the logical object so that the mirror cache is destaged to the at least one physical storage device, no portion of the logical object fragments being otherwise destaged; and

a remote protocol executable on the remote storage site ~~and adapted to control that~~ controls the cache sidefile to cache the multiple fragments as received and to destage the logical object to the storage on receipt of all fragments.

2. (Previously presented): The database system according to Claim 1 wherein: the main protocol includes information indicative of logical object fragment commencement and completion using a technique selected from among a group consisting of: (1) explicitly sending a start control message preceding the multiple fragments and an end control message concluding the multiple fragments, and (2) implicitly determining either a start control message or an end control message.
3. (Original): The database system according to Claim 1 wherein the main protocol further comprises:
an address translation process that translates a logical address to a list of physical addresses.
4. (Original): The database system according to Claim 1 wherein the main protocol further comprises:
an address translation process that resolves a virtual write address of the database application into a pick list of actual physical media writes associated with the logical object.
5. (Previously presented): The database system according to Claim 1 wherein the main protocol further comprises:
a process adapted to create a control message for communication to the remote protocol that instructs individual physical storage elements to operate on the multiple physical writes as a single object entity so that all or none is destaged to the storage.
6. (Canceled).
7. (Original): The database system according to Claim 1 wherein:
information is replicated from the main storage site to the remote storage site using a technique selected from among a group including: (1) synchronous data replication and (2) asynchronous data replication.

8. (Original): The database system according to Claim 1 wherein:
the logical object multiple fragments are controllably destaged in all-or-none fashion
to all devices in a consistency group.

9. (Currently Amended): An article of manufacture comprising:
a controller usable medium having a computable readable program code embodied
therein for executing in a database system that runs a database application for
mirroring a logical object in multiple fragments from a main storage site to a
remote storage site, the computable readable program code further
comprising:

a code configured to cause the controller to interface with the database
application that links and mirrors data between the main storage site
and the remote storage site, the remote storage site including a storage
and a cache sidefile divided into a plurality of array sidefile
recordsets; and

a code configured to cause the controller to create and deploy the logical
object in multiple fragments in combination with control information
indicative of logical object fragment commencement and completion
in the multiple fragment database application transfer; the control
information controlling the cache sidefile to cache the multiple
fragments as received and to destage the logical object to the storage
on receipt of all fragments;

a code adapted to receive an application request to write the logical object of
a specified length to a specified virtualized storage address;

a code adapted to convert the virtualized write address and resolving the
transfer length to designate at least one physical address in at least one
physical storage device for transferring the logical object in
fragments;

a code adapted to send a first control message to the at least one physical
storage device that delineates the start of a logical object that is to be
held in a remote mirror cache for destaging; and

a code adapted to send a second control message that delineates the end of the
logical object so that the mirror cache is destaged to the at least one

physical storage device, no portion of the logical object fragments being otherwise destaged.

10. (Previously presented): The article of manufacture according to Claim 9 wherein the computable readable program code further comprises:

a code adapted to create control information indicative of logical object fragment commencement and completion using a technique selected from among a group consisting of: (1) explicitly sending a start control message preceding the multiple fragments and an end control message concluding the multiple fragments, and (2) implicitly determining either a start control message or an end control message.

11. (Previously presented): The article of manufacture according to Claim 9 wherein the computable readable program code further comprises:

a code adapted to translate a logical address to a list of physical addresses.

12. (Previously presented): The article of manufacture according to Claim 9 wherein the computable readable program code further comprises:

a code adapted to resolve a virtual write address of the database application into a pick list of actual physical media writes associated with the logical object.

13. (Previously presented): The article of manufacture according to Claim 9 wherein the computable readable program code further comprises:

a code adapted to create a control message for communication to the remote protocol that instructs individual physical storage elements to operate on the multiple physical writes as a single object entity so that all or none is destaged to the storage.

14. (Canceled).

15. (Previously presented): The article of manufacture according to Claim 9 wherein the computable readable program code further comprises:

a code adapted to replicate information from the main storage site to the remote storage site using a technique selected from among a group including: (1) synchronous data replication and (2) asynchronous data replication.

16. (Previously presented): The article of manufacture according to Claim 9 wherein the computable readable program code further comprises:

a code adapted to controllably destage the logical object multiple fragments in all-or-none fashion to all devices in a consistency group.

17. (Currently Amended): An article of manufacture comprising:

a controller usable medium having a computable readable program code embodied therein for executing in a database system that runs a database application for mirroring a logical object in multiple fragments from a main storage site to a remote storage site, the computable readable program code further comprising:

a code executable at the remote storage site configured to cause the controller to receive the logical object in multiple fragment transfers in combination with control information indicative of logical object fragment commencement and completion;

a code executable at the remote storage site configured to cause the controller to control storage of the logical object multiple fragments in a cache sidefile divided into a plurality of array sidefile recordsets;

a code executable at the remote storage site configured to cause the controller to receive first and second control messages from a main storage site identifying respective start and end of the logical object; and

a code executable at the remote storage site configured to cause the controller to cache the multiple fragments as received and to destage the logical object to the storage on receipt of all fragments according to the first and second control messages.

18. (Previously presented): The article of manufacture according to Claim 17 wherein the computable readable program code further comprises:

a code adapted to determine logical object fragment commencement and completion using a technique selected from among a group consisting of: (1) receiving

explicitly identified starting and ending fragments, and (2) deriving either of the starting fragment and the ending fragment implicitly from received control information.

19. (Previously presented): The article of manufacture according to Claim 17 wherein the computable readable program code further comprises:
a code configured to cause the controller to track order of fragment updating between the main storage site and the remote storage site including updating of the sidefile recordsets.

20. (Previously presented): The article of manufacture according to Claim 17 wherein the computable readable program code further comprises:
a code adapted to replicate information from the main storage site to the remote storage site using a technique selected from among a group including: (1) synchronous data replication and (2) asynchronous data replication.

21. (Previously presented): The article of manufacture according to Claim 17 wherein the computable readable program code further comprises:
a code adapted to controllably destage the logical object multiple fragments in all-or-none fashion to all devices in a consistency group.

22. (Currently Amended): A storage element readable by a controller tangibly embodying a program of instructions executable by the controller to perform method acts for executing in a database system that runs a database application for mirroring a logical object in multiple fragments from a main storage site to a remote storage site, the method acts comprising:

receiving the logical object at the remote storage site in multiple fragment transfers in combination with control information indicative of logical object fragment commencement and completion;

controlling storage of the logical object multiple fragments at the remote storage site in a cache sidefile divided into a plurality of array sidefile recordsets;

receiving at the remote storage site first and second control messages from a main storage site identifying respective start and end of the logical object;

caching the multiple fragments at the remote storage site as received; and

destaging the logical object at the remote storage site to the storage on receipt of all fragments.

23. (Previously presented): The storage element according to Claim 22 wherein the method acts further comprise:

creating control information indicative of logical object fragment commencement and completion using a technique selected from among a group consisting of: (1) explicitly identifying starting and ending fragments, and (2) implicitly indicating either of the starting fragment and the ending fragment.

24. (Original): The storage element according to Claim 22 wherein the method acts further comprise:

resolving a virtual write address of the database application into a pick list of actual physical media writes associated with the logical object.

25. (Original): The storage element according to Claim 22 wherein the method acts further comprise:

creating a control message for communication to the remote protocol that instructs individual physical storage elements to operate on the multiple physical writes as a single object entity so that all or none is destaged to the storage.

26. (Currently Amended): The storage element according to Claim 22 wherein the method acts further comprise:

receiving at the main storage site an application request to write the logical object of a specified length to a specified virtualized storage address;
converting at the main storage site the virtualized write address and resolving the transfer length to designate at least one physical address in at least one physical storage device for transferring the logical object in fragments;
sending a first control message from the main storage site to the at least one physical storage device that delineates the start of a logical object that is to be held in a remote mirror cache for destaging; and
sending a second control message that delineates the end of the logical object so that the mirror cache is destaged to the at least one physical storage device, no portion of the logical object fragments being otherwise destaged.

27. (Original): The storage element according to Claim 22 wherein the method acts further comprise:

replicating information from the main storage site to the remote storage site using a technique selected from among a group including: (1) synchronous data replication and (2) asynchronous data replication.

28. (Currently Amended): A storage element readable by a controller tangibly embodying a program of instructions executable by the controller to perform method acts for executing in a database system that runs a database application for mirroring a logical object in multiple fragments from a main storage site to a remote storage site, the method acts comprising:

interfacing with the database application that links and mirrors data between the main storage site and the remote storage site, the remote storage site including a storage and a cache sidefile divided into a plurality of array sidefile recordsets;

deploying from the main storage site the logical object in multiple fragments in combination with control information indicative of logical object fragment commencement and completion in the multiple fragment database application transfer; the control information controlling the cache sidefile to cache the multiple fragments as received;

receiving an application request to write the logical object of a specified length to a specified virtualized storage address;

converting the virtualized write address and resolving the transfer length to designate at least one physical address in at least one physical storage device for transferring the logical object in fragments;

sending a first control message to the at least one physical storage device identifying the start of a logical object to be held in a remote mirror cache for destaging;

sending a second control message identifying the end of the logical object so that the mirror cache is destaged to the at least one physical storage device, no portion of the logical object fragments being otherwise destaged; and

destaging at the remote storage site the logical object to the storage on receipt of all fragments.

29. (Original): The storage element according to Claim 28 wherein the method acts further comprise:
determining logical object fragment commencement and completion using a technique selected from among a group including: (1) receiving explicitly identified starting and ending fragments, and (2) deriving either of the starting fragment and the ending fragment implicitly from received control information.
30. (Original): The storage element according to Claim 28 wherein the method acts further comprise:
tracking order of fragment updating between the main storage site and the remote storage site including updating of the sidefile recordsets.
31. (Original): The storage element according to Claim 28 wherein the method acts further comprise:
replicating information from the main storage site to the remote storage site using a technique selected from among a group including: (1) synchronous data replication and (2) asynchronous data replication.
32. (Original): The storage element according to Claim 28 wherein the method acts further comprise:
controllably destaging the logical object multiple fragments in all-or-none fashion to all devices in a consistency group.